

Serial No.: 10/620,080

RECEIVED
CENTRAL FAX CENTER

JAN 16 2007

REMARKS

Claims 1-10 and 23-32 are pending in the present application. In view of the following remarks, favorable reconsideration of this case is respectfully requested.

Claims 1, 2, 5-10, 23, 24, and 27-32 are rejected under 35 U.S.C. 102(b) as being unpatentable over United States Patent No. 5,615,352 to Jacobson et al. (hereinafter referred to as Jacobson). Applicants respectfully traverse.

Claim 1 relates to a method for data distribution that includes, *inter alia*, distributing logical addresses among an initial set of storage devices so as provide a balanced access to the devices, transferring the data to the storage devices in accordance with the logical addresses, and adding an additional storage device to the initial set, thus forming an extended set of the storage devices comprising the initial set and the additional storage device. The method of claim 1 also includes *redistributing the logical addresses among the storage devices in the extended set so as to cause a portion of the logical addresses to be transferred from the storage devices in the initial set to the additional storage device*. In claim 1, the redistributing is performed *while maintaining the balanced access and while maintaining the same logical addresses for the logical addresses in the initial set of storage devices that are not transferred to the additional storage device*.

The Examiner asserts that Jacobson discloses all of the features of claim 1. However, Jacobson does not disclose or suggest redistributing the logical addresses among the storage devices in the extended set so as to cause a portion of the logical addresses to be transferred from the storage devices in the initial set *to the additional storage device*. The Examiner cites Jacobson; col. 3, lines 5-13, as alleged disclosure of this feature. However, the cited section is presented below within the context of the entire paragraph:

Serial No.: 10/620,080

In this case, when one or more storage disks are added to the hierarchic disk array, data stored in *a selected RAID area is moved to another area of equal or greater storage capacity*. The mapping of first to second virtual storage spaces is updated to reflect the data movement. The physical storage space corresponding to the selected RAID area is then remapped into an expanded RAID area which spans across all storage disks, including the new disks. This process is repeated until all RAID areas in the hierarchic disk array have been expanded to include regions on the new storage disks. The mapping of the second virtual storage space to first is continuously modified to reflect the expansion of the RAID areas within the virtual storage space.

(Jacobson; col. 3, lines 1-13; emphasis added). Jacobson apparently discloses adding disk space to a virtual stripe configuration of memory, and apparently mentions moving a selected RAID area to another area of storage. However, this does not disclose or suggest redistributing the logical addresses to cause a portion of the logical addresses to be transferred from the storage devices in the initial set *to the additional storage device*. Jacobson apparently adds disk space and redistributes the addresses (though not necessarily in a load-balanced manner) so that data from one stripe may be moved to another portion of the physical storage space. However, this does not disclose or suggest that *data is moved into another portion of the new storage areas in the physical storage space*. Therefore, for at least this reason claim 1 is allowable.

Additionally, Jacobson does not disclose or suggest that the redistributing is performed *while maintaining the balanced access*. The Examiner cites the same section above in support of this feature. However, the cited section gives no indication of *maintaining balanced access* after expanding storage capacity. Balanced access is discussed and defined in the specification at least at page 17, lines 10-17, which states:

The procedures allocate the logical stripes to devices B_n so that balanced access to the devices is maintained, where balanced access assumes that taken over approximately $10,000 \times N$ transactions with devices B_n , the fraction of capacities of devices B_n used are equal to within approximately 1%, where N is the

Serial No.: 10/620,080

number of devices B_n , the values being based on a Bernoulli distribution.

Jacobson apparently mentions moving a selected RAID area to an area of equal or greater capacity, and apparently discusses remapping the virtual storage area after expanding the storage capacity. However, it is respectfully submitted that Jacobson does not suggest that the storage device therein has balanced access after adding storage capacity, and therefore for at least this additional reason claim 1 is allowable.

Furthermore, Jacobson does not disclose or suggest that the redistributing is performed *while maintaining the same logical addresses for the logical addresses in the initial set of storage devices that are not transferred to the additional storage device*. The Examiner cites Jacobson; col. 2, lines 9-25, as alleged disclosure of this feature. The cited section is presented below within the context of the entire paragraph:

The present invention provides methods for enlarging storage capacity in a data storage system by adding more storage disks, yet still maintaining data availability. According to one method, the physical storage space of the disk array is configured into multiple stripes for storing a predetermined amount of data. The stripes extend across multiple storage disks in the disk array, and are made up of one or more equal sized segments from each storage disk in the disk array. When more storage disks are added, *data from one stripe is moved to a another portion of the physical storage space*. The physical storage space containing the stripe is then reconfigured into an expanded stripe for storing data. This expanded stripe spans across all storage disks, including the new additional storage disks. The expanded stripe is then ready to receive new data. This process is continued stripe-by-stripe until all stripes have been configured to include the new disks. While the stripe is being expanded, the data storage system does not allocate any virtual blocks within the selected stripe.

(Jacobson; col. 2, lines 6-24; emphasis added). There is no limitation presented in Jacobson that data that is in a stripe prior to expansion is *not moved* to another storage area of the stripe, which is not part of the expanded storage area, after expansion. In contrast, Applicant's invention

Serial No.: 10/620,080

recites moving data *only* into the added space, and never moving data within the original storage space after expanding the storage area. Jacobson moves data within the storage space, but not necessarily *only* to the added space.

The present invention discusses how to do add additional storage with a minimal transfer of data, and while maintaining load balance. The present specification discloses transferring only the data which will reach the new disk after the additional storage space is added and the load is balanced, *without moving other data in the system*. Jacobson does not disclose this feature, but instead *rearranges the entire address space*. The present invention avoids rearranging the entire address space by moving some addresses to the new space, while leaving other addresses in their original position. The present invention therefore provides a more efficient method for reassigning logical addresses, and reduces the opportunity for error, by preserving load balance while minimizing the migration of data in the new distribution. There is no disclosure in any of the cited sections of Jacobson of performing *load balancing* and *redistributing, while maintaining the same logical addresses for the logical addresses in the initial set of storage devices that are not transferred to the additional storage device*. Therefore, for at least this reason claim 1 is allowable over Jacobson.

Claims 1, 2, and 5-10 depend from claim 1 and therefore these claims are allowable for at least the same reasons claim 1 is allowable.

Amended claim 23 includes a feature similar to that discussed above in regard to amended claim 1, and therefore, for at least the same reasons claim 1 is allowable, claim 23 is also allowable.

Claims 24, and 27-32 depend from claim 23 and therefore these claims are allowable for at least the same reasons claim 23 is allowable.

Serial No.: 10/620,080

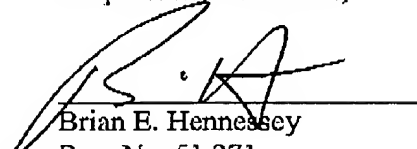
Claims 3, 4, 25, and 26 are rejected under 35 U.S.C. 103(a) as being unpatentable over Jacobson and further in view of "Consistent Hashing and Random Trees: Distributed Caching Protocols for Relieving Hot Spots on the Worldwide Web," by Karger et al., in the Proceedings of the 29th ASM Symposium on Theory of Computing, Pages 654-663 (hereinafter referred to as Karger). Applicants respectfully traverse.

The addition of Karger fails to cure the critical deficiency discussed above as regards Jacobson as applied against claims 1 and 23, et al. Therefore, since claims 3 and 4 depend from claim 1, these claims are allowable for at least the same reasons claim 1 is allowable. Similarly, since claims 25 and 26 depend from claim 23, these claims are allowable for at least the same reasons claim 23 is allowable.

In view of the remarks set forth above, this application is believed to be in condition for allowance which action is respectfully requested. However, if for any reason the Examiner should consider this application not to be in condition for allowance, the Examiner is respectfully requested to telephone the undersigned attorney at the number listed below prior to issuing a further Action.

Any fee due with this paper may be charged on Deposit Account 50-1290.

Respectfully submitted,



Brian E. Hennessey
Reg. No. 51,271

CUSTOMER NO.: 026304
Phone No.: (212) 940-6311
Fax No.: (212) 940-8986/7
DOCKET NO.: ASSIA 20.502